

What is claimed is:

1 Claim 1. A multiprocessor computer system, comprising:
2 a plurality of processing nodes and a plurality of
3 dynamic cache coherency regions using caches associated with
4 said processing nodes, and having
5 cache controller logic in said processing nodes
6 controlling movement of software initiated movement of
7 software processes between said plurality of cache coherency
8 regions without requiring a selective purging of cache
9 contents in one or more of said processing nodes.

1 Claim 2. The multiprocessor computer system according to
2 claim 1, including supervisor software which enables said
3 cache controller logic in a processing node to be sure upon
4 an incoming storage request for a storage address that no
5 copy of the requested storage address exists outside that
6 processor's current coherency region, as specified by the
7 current coherency region mode, whenever a cache entry for a
8 requested storage address is found to exist on any cache in
9 said processing nodes that contains a processor that
10 initiated said incoming storage request.

1 Claim 3. The multiprocessor computer system according to
2 claim 2, wherein said supervisor software creates a unique
3 Coherency Region ID for each process that has its own
4 coherency region.

1 Claim 4. The multiprocessor computer system according to
2 claim 2, wherein supervisor software creates a table for
3 each processing node in the system which has an entry for
4 every Coherency Region ID that is currently allowed to be
5 dispatched on said processing node.

1 Claim 5. The multiprocessor computer system according to
2 claim 2, wherein said supervisor software creates a unique
3 Coherency Region ID for each process that has its own
4 coherency region and one or more coherency mode bits for
5 each processor in the multiprocessor computer system, and
6 said coherency mode bits and coherency region ID associated
7 with a processor are sent together with each storage
8 transaction that is initiated by that processor when the
9 transaction is transmitted for communication to another
10 processor of said multiprocessor computer system.

1 Claim 6. The multiprocessor computer system according to
2 claim 2, wherein said supervisor software creates a unique
3 Coherency Region ID for each process that has its own
4 coherency region and one or more coherency mode bits for
5 each processor in the multiprocessor system to
6 to a node controller for a processing node.

1 Claim 7. The multiprocessor computer system according to
2 claim 2, wherein said supervisor software creates a unique
3 Coherency Region ID for each process that has its own
4 coherency region and one or more coherency mode bits for
5 each processor in the multiprocessor computer system, and
6 wherein said mode bits associated with each transaction are
7 used to determine which caches must participate in any
8 storage transactions that they receive from any of the
9 processors of said multiprocessor computer system.

1 Claim 8. The multiprocessor computer system according to
2 claim 2, wherein said supervisor software creates a unique
3 Coherency Region ID for each process that has its own
4 coherency region and one or more coherency mode bits for
5 each processor in the multiprocessor computer system and

6 enables multiple cache coherency regions to operate without
7 the use of cache purges during some operations which move
8 software processes between coherency regions.

1 Claim 9. The multiprocessor computer system according to
2 claim 8, wherein said supervisor software moves a software
3 process out of one coherency region that is no longer going
4 to be used by said software process and into another
5 software process that has been created to cover the same
6 address space as the first but which will include a new set
of processing nodes.

1
2 Claim 10. The multiprocessor computer system according to
3 claim 2, wherein said supervisor software creates a unique
4 Coherency Region ID for each process that has its own
5 coherency region and moves a software process from one
6 coherency region encompassing one set of processing nodes to
7 another coherency region encompassing another set of
8 processing nodes without requiring cache purges of caches in
any of the processing nodes.

1
2 Claim 11. The multiprocessor computer system according to
3 claim 10, wherein if said another coherency region contains
4 fewer hardware processing nodes than the original coherency
5 region, the size of the coherency region for said processing
nodes has been effectively been reduced.

1
2 Claim 12. The multiprocessor computer system according to
3 claim 1, wherein said multiprocessor computer system having
4 a plurality of said processsing nodes uses a table of active
5 coherency region information associated with each processing
6 node to determine when to alter the processing nodes' cache
state transitions.

1

2 Claim 13. The multiprocessor computer system according to
3 claim 12, wherein a supervisor program initializes said
4 tables associated with each processing node and an entry in
5 said table is made for each coherency region that the
supervisor program intends to use on that processing node.

1

2 Claim 14. The multiprocessor computer system according to
3 claim 1, wherein a supervisor program assigns a unique
4 coherency region ID for each coherency region which the
5 supervisor can associate with all software processes that
6 access storage addresses that are encompassed by the
coherency region.

1

2 Claim 15. The multiprocessor computer system according to
3 claim 1, wherein processing nodes are able to identify
4 incoming storage requests which target lines that are no
5 longer part of the address space of any software process
6 that is currently enabled by the supervisor software to be
dispatched on the node.

1

2 Claim 16. The multiprocessor computer system according to
3 claim 1, wherein processing nodes are able to identify
4 incoming storage requests which target lines that are no
5 longer part of the address space of any software process
6 that is currently enabled by the supervisor software to be
7 dispatched on the node to thereby identify cache lines that
8 are no longer actively used by any software processes on
9 that processing node and to change the cache entries for
10 that processing node to invalid in response to a storage
request from outside the coherency region.

1

2

3 Claim 17. The multiprocessor computer system according to
3 claim 1, wherein processing nodes are able to identify
4 incoming storage requests which target lines that are no
5 longer part of the address space of any software process
6 that is currently enabled by the supervisor software to be
7 dispatched on the node and allows all of the caches in said
8 multiprocessor computer system to continue processing
coherency transactions while the coherency boundaries for a
1 software process are effectively changed.

2

3 Claim 18. The multiprocessor computer system according to
4 claim 1, wherein processing nodes are able to identify
5 incoming storage requests which target lines that are no
6 longer part of the address space of any software process
7 that is currently enabled by the supervisor software to be
8 dispatched on the node such that cache lines belonging to a
9 software process that is no longer actively being dispatched
on a given processing node are identified and invalidated
1 thereby enabling their reuse.

2

3 Claim 19. The multiprocessor computer system according to
4 claim 1, wherein a supervisor software uses processor state
5 information to determine which caches in the multiprocessor
6 computer system are required to examine a coherency
transaction produced by a single originating processor's
1 incoming storage request.

2

3 Claim 20. The multiprocessor computer system according to
4 claim 19, wherein a processing node of the multiprocessor
5 computer system has has dynamic coherency boundaries such
6 that the multiprocessor computer system uses only a subset
7 of the total processors in a system for a single workload at
any specific point in time and can optimize the cache

8 coherency as the supervisor software expands and contracts
9 the number of processors which are being used to run any
single workload.

1

2 Claim 21. The multiprocessor computer system according to
3 claim 20, wherein multiple instances of processing nodes can
4 be connected with a second level controller to create a
large multiprocessor system.

1

2 Claim 22. The multiprocessor computer system according to
3 claim 21, wherein said supervisor software creates a unique
4 Coherency Region ID for each process that has its own
5 coherency region and one or more coherency mode bits for
6 each processor in the multiprocessor computer system and
7 enables multiple cache coherency regions to operate without
8 the use of cache purges during some operations which move
9 software processes between coherency regions and a node
10 controller uses said mode bits to determine which processors
11 must receive any given transaction that is received by the
node controller.

1

2 Claim 23. The multiprocessor computer system according to
3 claim 22, wherein a second level controller uses the mode
4 bits to determine which processing nodes must receive any
5 given transaction that is received by the second level
controller.

1

2 Claim 24. The multiprocessor computer system according to
3 claim 21, wherein said supervision software uses logical
4 partitions which are mapped to allowable physical processors
5 and a distinct cache coherency region can be defined for
each partition using a hypervisor.

1 Claim 25. The multiprocessor computer system according to
2 claim 2, wherein said coherency region ID is used to perform
3 the function of a cache coherency mode and a node controller
4 determines which physical processing nodes are associated
5 with specific coherency region Ids.

1 Claim 26. The multiprocessor computer system according to
2 claim 3, wherein any incoming storage request which misses
3 all of the caches in an originator's coherency region is
4 then sent on to all processing nodes in the entire system,
5 regardless of the setting of said mode bits.

1 Claim 27. The multiprocessor computer system according to
2 claim 3, wherein any incoming storage request which
3 Requests which hit in an originator's coherency region but
4 which do not have a correct cache state do not need to be
5 sent outside the coherency region.

1 Claim 27. The multiprocessor computer system according to
2 claim 3, wherein a cache line cannot be marked as shared in
3 two separate coherency regions, and when said supervisor
4 program is moving a coherency region from one set of
5 processing nodes to another set of processing nodes it is
6 effectively leaving behind cache entries for the coherency
7 region on the old nodes and ensures that these old cache
8 entries will be seen by incoming stroage requests
9 originating from the new processing nodes and that cache
10 entries for the same main storage addresses will not be
11 established in the new processing nodes until the old
12 entries are invalidated.

1 Claim 28. A method for use in a multiprocessor computer
2 system, comprising the steps of:

3 moving software processes between a plurality of
4 cache coherency regions for caches associated with a
5 plurality of processing nodes of said multiprocessor
6 computer system without requiring a selective purging of
7 cache contents in one or more of said processing nodes,
8 after supervisor software creates a unique Coherency Region
9 ID for each process that has its own coherency region, and
10 said supervisor software creates a table for each
11 processing node in the multiprocessor computer system which
12 has an entry for every Coherency Region ID that is currently
13 allowed to be dispatched on said processing node.

1 Claim 29. The multiprocessor computer system according to
2 claim 28, wherein said coherency mode bits and coherency
3 region ID associated with a processor are sent together with
4 each storage transaction that is initiated by that processor
5 with a requested storage address when the transaction is
6 transmitted for communication to another processor of said
7 multiprocessor computer system.

1 Claim 30. The multiprocessor computer system according to
2 claim 28, including a step of enabling with supervisor
3 software the multiprocessor computer system's cache
4 controller logic in a processing node to be sure that, upon
5 receipt at said processing node that an incoming storage
6 request for a storage address, no copy of the requested
7 storage address exists outside said processing node's
8 current coherency region, as specified by a current
9 coherency region mode, whenever a cache entry for a
10 requested storage address is found to exist on any cache in
11 any of said multiprocessor computer system's processing
12 nodes that contains a processor that initiated said incoming
13 storage request.